

Calibration of structural and reduced-form recovery models

Alexander F. R. Koivusalo

Danske Capital, Copenhagen, Denmark

Faculty of Physics, University of Duisburg-Essen, Germany

Rudi Schäfer

Faculty of Physics, University of Duisburg-Essen, Germany

Abstract

In recent years research on credit risk modelling has mainly focused on default probabilities. Recovery rates are usually modelled independently, quite often they are even assumed constant. Then, however, the structural connection between recovery rates and default probabilities is lost and the tails of the loss distribution can be underestimated considerably. The problem of underestimating tail losses becomes even more severe, when calibration issues are taken into account. To demonstrate this we choose a Merton-type structural model as our reference system. Diffusion and jump-diffusion are considered as underlying processes. We run Monte Carlo simulations of this model and calibrate different recovery models to the simulation data. For simplicity, we take the default probabilities directly from the simulation data. We compare a reduced-form model for recoveries with a constant recovery approach. In addition, we consider a functional dependence between recovery rates and default probabilities. This dependence can be derived analytically for the diffusion case. We find that the constant recovery approach drastically and systematically underestimates the tail of the loss distribution. The reduced-form recovery model shows better results, when all simulation data is used for calibration. However, if we restrict the simulation data used for calibration, the results for the reduced-form model deteriorate. We find the most reliable and stable results, when we make use of the functional dependence between recovery rates and default probabilities.

Keywords: Credit risk, Loss distribution, Reduced-form models, Structural models, Value at Risk, Expected Tail Loss, Stochastic processes

JEL: C15, G21, G24, G28, G33

1. Introduction

To correctly assess and control exposure to credit risk is a vital question for financial institutions. It is also a crucial problem for banking regulation. There are two conceptually different approaches to credit risk modelling: structural and reduced-form approaches. The structural models go back to Black and Scholes (1973) and Merton (1974). The Merton model is based on the assumption that a company has a certain amount of zero-coupon debt which becomes due at a fixed maturity date. The market value of the company is modelled as a stochastic process. A possible default and the associated recovery rate are determined directly from this market value at maturity. In the reduced-form approach default probabilities and recovery rates are described independently by stochastic models. The aim is to describe the dependence of these quantities on common covariates or risk factors. For some well known reduced-form model approaches see, e.g., Jarrow and Turnbull (1995), Jarrow et al. (1997), Duffie and Singleton (1999), Hull and White (2000) and Schönbucher (2003). First Passage Models constitute a third approach which is usually regarded as structural, but is better described as a mixed or pseudo-structural approach. First Passage Models were first introduced by Black and Cox (1976). As in the Merton model, the market value of a company is modelled as a stochastic process. Default occurs as soon as the market value falls below a certain threshold. In contrast to the Merton model, default can occur at any time. In this approach, the recovery rate is not determined by the underlying process for the market value. Instead, recovery rates are modelled independently, for example, by a reduced-form approach (see e.g., Asvanunt and Staal (2009a,b)). In some cases recovery rates are even assumed constant, for instance, in Giesecke (2004).

The independent modelling of default and recovery rates can lead to a serious underestimation of large losses. This situation is particularly troublesome when considering calibration to limited default and recovery data. One way to address the scarcity issue of default and recovery data is to calibrate to market data of credit contracts and derivatives instead. This allows for a self-consistent description of market prices, but may misrepresent the actual credit risk. As witnessed in the recent financial crisis, the market as a whole may be completely wrong in its pricing consensus. Therefore it is crucial for the risk assessment of a credit portfolio to take into account historical data of defaults and recoveries.

The aim of our study is to discuss general calibration issues that arise from independent modelling of default and recovery rates. To this end we use a structural model as our reference system. While the original Merton model contains only a diffusion process, we also consider a jump-diffusion process for the evolution of the market value. For related works on jump-diffusion processes in credit risk modelling see, e.g., Zhou (2001), Schäfer et al. (2007) and Kiesel and Scherer (2007). We run Monte Carlo simulations of this model and calibrate different recovery models to the simulation data. For simplicity, we take the default probabilities directly from the simulation data. We compare a reduced-form model for recoveries with a constant recovery approach. In addition we consider a functional dependence between recovery and default rates. This dependence can be derived analytically for the diffusion case. We find that the constant recovery approach drastically and systematically underestimates the tail of the loss distribution. The reduced-form recovery model

shows better results, when all simulation data is used for calibration. However, if we restrict the calibration data, the results for the reduced-form model deteriorate. We find the most reliable and stable results, when we make use of the functional dependence between recovery and default rates.

The paper is organized as follows. In Section 2 we discuss the structural model used as our reference system. In Section 3 we describe the recovery rate models which will be calibrated to the Monte Carlo data. Issues that arise in model calibration are discussed in Section 4. In Section 5 we compare the different recovery models with respect to three central questions: To what extent do the models preserve the dependence of recovery and default rates? How robust are the models with respect to calibration issues? And do they provide accurate estimates for the risk of large portfolio losses? We summarize our findings and conclude in Section 6.

2. Structural model as reference system

In the Merton model, defaults and recoveries are directly determined by an underlying market value at maturity. Hence, stochastic modelling of the market value $V_k(t)$ of a company allows to assess its credit risk. The model assumes that a company k has a certain amount of zero-coupon debt with a face value F_k . The debt will become due at maturity time T . The company defaults if the value of its assets at time T is less than the face value, i.e., if $V_k(T) < F_k$. The recovery rate is then $R_k = V_k(T)/F_k$ and the normalized loss given default is

$$L_k^* = 1 - R_k = \frac{F_k - V_k(T)}{F_k} . \quad (1)$$

We denote the loss given default with an asterisk to distinguish it from the loss including non-default events. The individual loss can be expressed as

$$L_k = \left(1 - \frac{V_k(T)}{F_k}\right) \Theta \left(1 - \frac{V_k(T)}{F_k}\right) , \quad (2)$$

where Θ is the Heaviside function.

A. Model setup

We model the time evolution of the market value of a single company k by a stochastic differential equation of the form

$$\frac{dV_k}{V_k} = \mu dt + \sqrt{c} \sigma dW_m + dJ_m + \sqrt{1-c} \sigma dW_k + dJ_k . \quad (3)$$

This is a correlated jump-diffusion process with a deterministic term μdt and a linearly correlated diffusion. The Wiener processes dW_k and dW_m describe the idiosyncratic and the market fluctuations, respectively. For simplicity, we choose the drift μ , the volatility σ and the correlation coefficient c as constant parameters, which are the same for all companies. The jump terms are not contained in Merton's original model; they ensure that the default

probability does not go to zero as the time to maturity becomes very short. We include two jump terms in the stochastic process, dJ_k for idiosyncratic jumps and dJ_m for jumps which affect the entire market. As in Schäfer et al. (2007), the jumps are modelled by a *Poisson process* with intensity λ . We recall that in such a process the probability function for the event to occur n times between zero and the time t is given by

$$p_n^{\text{Poisson}}(t) = \frac{(\lambda t)^n}{n!} \exp(-\lambda t) . \quad (4)$$

The size Λ of the jump, measured in units of the current asset value $V(t)$, is a random variable with a distribution which we have to specify. Jumps can be positive or negative. The largest possible negative jump is 100% of the current asset value. Based on this information, a possible distribution of the jump size Λ is a shifted lognormal distribution, $\Lambda + 1 \sim \text{LN}(\mu_J, \sigma_J)$, with mean μ_J and standard deviation σ_J .

Without the jump term, the distribution of the asset price $V_k(t)$ is log-normal. The jumps render the tails of the asset price distribution fatter. Fat tails are empirically observed Mantegna and Stanley (2000). As this clearly affects the loss distribution, we find it important to include such jumps. The parameters of the jump process can be adjusted in order to match the tail behavior of a given empirical time series of the asset value. We use the same parameters for idiosyncratic and market wide jumps.

B. Monte Carlo simulation

In the Monte Carlo simulations we consider the stochastic process in Equation (3) for discrete time increments $\Delta t = T/N$, where the time to maturity T is divided into N steps. The market value of company k at maturity is then given by

$$V_k(T) = V_k(0) \prod_{t=1}^N \left(1 + \mu \Delta t + \sqrt{c} \sigma \eta_{m,t} \sqrt{\Delta t} + dJ_{m,t} + \sqrt{1-c} \sigma \varepsilon_{k,t} \sqrt{\Delta t} + dJ_{k,t} \right) . \quad (5)$$

The random variables $\eta_{m,t}$ and $\varepsilon_{k,t}$ are independent and are drawn from a normal distribution. We consider a homogeneous portfolio of size K with the same parameters for each asset process, and with the same face value, $F_k = F$, and initial market value, $V_k(0) = V_0$. The simulation is run with an inner loop and an outer loop. In the inner loop we simulate $K = 500$ different realizations of $\varepsilon_{k,t}$ and $dJ_{k,t}$ for a single realization of the market terms $\eta_{m,t}$ and $dJ_{m,t}$ with $t = 1, \dots, N$. The inner loop can be interpreted as a homogeneous portfolio of size K , or simply as an average over the idiosyncratic part of the process. In each run of the inner loop, we calculate the market return X_m , the number of defaults $N_D(X_m)$ and the expected recovery rate $\langle R(X_m) \rangle$. The market return X_m is defined as the average return at maturity,

$$X_m = \frac{1}{K} \sum_{k=1}^K \left(\frac{V_k(T)}{V_k(0)} - 1 \right) . \quad (6)$$

For sufficiently large K the idiosyncratic terms average out and the market return X_m is solely defined by the realizations of $\eta_{m,t}$ and $dJ_{m,t}$. This is why we use the market return as

a parameter for the other observables. The number of defaults $N_D(X_m)$ simply counts how many times the condition $V_k(T) < F$ is fulfilled. We can estimate the default probability as

$$P_D(X_m) \approx N_D(X_m)/K . \quad (7)$$

We obtain the portfolio loss as the average of individual losses in Equation (2),

$$\langle L(X_m) \rangle = \frac{1}{K} \sum_{k=1}^K L_k . \quad (8)$$

Using the relation

$$\langle L(X_m) \rangle = P_D(X_m) (1 - \langle R(X_m) \rangle) \quad (9)$$

we can estimate the expected recovery rate as

$$\langle R(X_m) \rangle = 1 - \frac{\langle L(X_m) \rangle}{P_D(X_m)} \approx 1 - \frac{K \langle L(X_m) \rangle}{N_D(X_m)} . \quad (10)$$

Here, we assume that the number of defaults is strictly non-zero, which is justified for large K . The outer loop runs over 10^6 realizations of the market terms, where we obtain different values for the market return X_m and, consequently, the number of defaults $N_D(X_m)$ and the expected recovery rate $\langle R(X_m) \rangle$.

The Monte Carlo simulations are carried out for diffusion and jump diffusion. In this paper we do not discuss the parameter dependence of the models or aim at calibrating them to a given portfolio. Instead, we only present the results for a single set of parameters with economically sensible values. This suffices to make our main argument about different recovery model approaches. As correlation coefficient we choose $c = 0.5$. The parameters for the diffusion process are $\mu = 0.05$ and $\sigma = 0.15$. The additional parameters for the jump terms are $\lambda = 0.005$, $\mu_j = 0.4$ and $\sigma_j = 0.3$. The initial market value is set to $V_0 = 100$, the face value of the zero-coupon bonds is $F = 75$ with maturity time $T = 1$.

3. Recovery rate models

In our study we want to isolate the influence of the different recovery models on various risk measures. Therefore we will not employ a stochastic description of default probabilities. Instead we consider the default data from the simulations directly. We compare three different recovery rate models: a reduced-form model, constant recoveries and structural recoveries.

A. Reduced-form approach

Reduced-form models are stochastic models for default probabilities and recovery rates. Their aim is to describe the dependence of these quantities on common (macroeconomic) covariates or risk factors. In our structural reference model we use a one-factor model with a constant correlation of all companies to the market. Therefore we use the market return, i.e., the average return of all assets at maturity, as a covariate.

Default events are commonly modelled by a doubly stochastic Poisson process or *Cox process*. As mentioned above, we will not consider a stochastic description of default probabilities in this paper. Instead we consider the default data from the simulations directly.

Reduced-form recovery rate models describe a monotonous dependence on covariates. In our case, for instance, we expect a lower recovery rate for lower market returns. Since the natural domain for the recovery rate is the interval $[0, 1]$, a generalized linear model is typically used. In this approach a linear dependence on the covariates is transformed by a link function to the desired domain. Recovery rate models commonly use either a logit or a probit link function. A logit link function is used, eg, in Schönbucher (2003), a probit link function is, eg, in Andersen and Sidenius (2005). As pointed out by Chava et al. (2008), both link functions lead to very similar results. In our study we focus on the probit model; then the reduced-form recovery rate model reads

$$\langle R(X_m) \rangle = \Phi(-\gamma X_m - \delta) \quad (11)$$

with the parameters γ and δ . The qualitative insights presented in this study do not depend on the choice of the link function.

B. Constant recovery rate

In the simplest reduced-form ansatz the recovery rate is assumed constant,

$$\langle R(X_m) \rangle = \langle R \rangle = \text{const} . \quad (12)$$

Although it is obvious that this crude model neglects important structural information and therefore cannot be accurate, it is still widely used in research and in practice. The only valid statements to make with a constant recovery rate involve setting $\langle R_k \rangle$ to zero and thereby estimating worst case losses.

C. Structural recovery rate

In addition to the reduced-form approach, we consider a functional dependence of recovery rates and default probabilities,

$$\langle R(P_D) \rangle = \frac{1}{P_D} \exp \left(-B \Phi^{-1}(P_D) + \frac{1}{2} B^2 \right) \Phi \left(\Phi^{-1}(P_D) - B \right) . \quad (13)$$

This result has been derived by Schäfer and Koivusalo (2011) in the framework of the Merton model for a correlated diffusion process, but it is also a very good approximation for other processes. We call this the structural recovery rate model. The relation depends only on a single parameter B . In the case of the correlated diffusion, B is determined by the parameters of this process.

4. Model calibration

We will use the data obtained in our Monte Carlo simulations for model calibration. Since we do not employ a model for default probabilities, these will be directly calculated from the number of defaults found in the simulation,

$$P_D(X_m) \approx \frac{N_D(X_m)}{K} . \quad (14)$$

The three different recovery models will be calibrated to the average recovery rates $\langle R(X_m) \rangle$ found in the simulation.

A. Reduced-form recovery rate

We apply Φ^{-1} to both sides of Equation (11) and obtain

$$b(X_m) = \Phi^{-1}(\langle R(X_m) \rangle) = -\gamma X_m - \delta . \quad (15)$$

According to our reduced-form model we expect $b(X_m)$ to show a linear dependence. In Figure 1 we present the Monte Carlo data of $b(X_m)$ for diffusion and jump-diffusion. In addition to the raw data, we also show average values for small intervals of X_m . The dependence on X_m is well defined and rather linear for large negative market returns. This is true both for diffusion and for jump-diffusion. For market returns greater than zero, the dependence on X_m is only visible on average. In single runs of the simulation, the recovery rates, and thus also $b(X_m)$, can deviate considerably from the average values. The reason for this is rather obvious: positive market returns make defaults less likely for all individual companies. Thus, defaults and recoveries are mostly determined by the idiosyncratic parts of the processes. The model assumption of a linear behavior of $b(X_m)$ is well justified for large negative market returns X_m . However, we observe deviations of this linear behavior starting already around $X_m = -0.2$. These deviations are most visible in the jump-diffusion case, where we even find a negative slope for $X_m \geq 0$. From these findings we can already deduce a qualitative picture of the calibration difficulties that may arise. The overall behavior of recovery rates is not captured by the reduced-form model. If we calibrate only to data for large negative market returns, we will overestimate recovery rates for moderate market returns. And if we only have data available for calibration which corresponds to rather moderate market returns, we would overestimate recovery rates for large negative market returns. In both cases, we overestimate recovery rates in some situations, and thus underestimate the corresponding losses. Here we choose 0 as an upper threshold for X_m and study the calibration results in dependence of a lower threshold. The motivation for this is the typical abundance of data for moderate market returns, while large negative market returns occur only scarcely. We calibrate to the local average values in order to give equal weights to different values of X_m in the regression.

Another problem becomes obvious when we compare the diffusion and the jump-diffusion case: The quality of the reduced-form recovery model depends on the underlying process. In other words, a model that works well for one set of credit contracts might not work so well for another. Non-stationarity in financial time series may also cause the quality of a reduced-form recovery model to degrade over time.

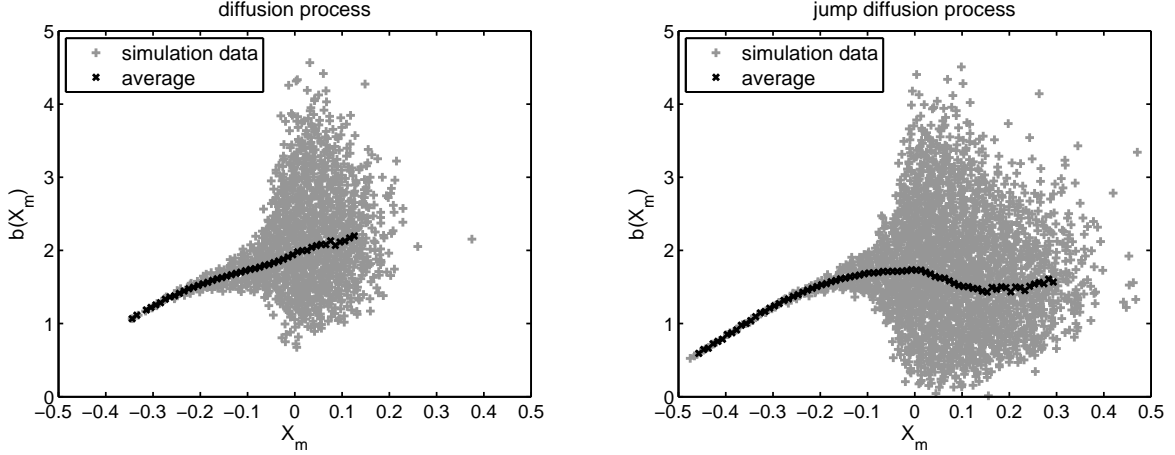


Figure 1: Regression data $b(X_m) = \Phi^{-1}(\langle R(X_m) \rangle)$ for the reduced-form recovery model. The model assumes a linear dependence, which is justified for large negative market returns X_m . Results are shown for diffusion (left), and for jump-diffusion (right).

B. Constant recovery rate

The constant recovery rate is determined as the average value of $\langle R(X_m) \rangle$ for a given range of market returns X_m . The reasons for choosing an upper and a lower threshold on X_m have been discussed above. The upper threshold $X_m < 0$ addresses the bias issue of the estimation at least to some degree.

C. Structural recovery rate

In our structural recovery approach, we do not fit a covariate dependence of the recovery rate explicitly. Instead, we focus on how recovery rates depend on default probabilities. The functional dependence of recovery rate on default probability depends only on a single parameter B .

$$\langle R_k(P_D) \rangle = \frac{1}{P_D} \exp \left(-B \Phi^{-1}(P_D) + \frac{1}{2} B^2 \right) \Phi(\Phi^{-1}(P_D) - B) . \quad (16)$$

In the Merton model for a correlated diffusion process the parameter B is given by the parameters of this process,

$$B = \sqrt{(1 - c)\sigma^2 T} . \quad (17)$$

We use a least square method to estimate B from the simulation data for $P_D(X_m)$ and $\langle R_k(X_m) \rangle$. Again, the data are filtered with respect to an upper and a lower threshold on the market return X_m . The convergence of the regression for B is strong. As already pointed out in Schäfer and Koivusalo (2011), this functional dependence does not critically depend on the underlying process. In fact, it has been shown to work well for diffusion, jump diffusion and GARCH(1,1). Therefore we do not expect a critical dependence of the calibration results on the range for X_m . We can further improve the stability of the calibration if we regress

the data of average losses directly. Inserting the functional dependence (13) into Equation (A.7) yields

$$\langle L(P_D) \rangle = P_D - \exp \left(-B \Phi^{-1}(P_D) + \frac{1}{2} B^2 \right) \Phi \left(\Phi^{-1}(P_D) - B \right) . \quad (18)$$

Now the strongly fluctuating results for small default probabilities are suppressed in the regression, see Figure 2. This leads to an improvement in calibration stability.

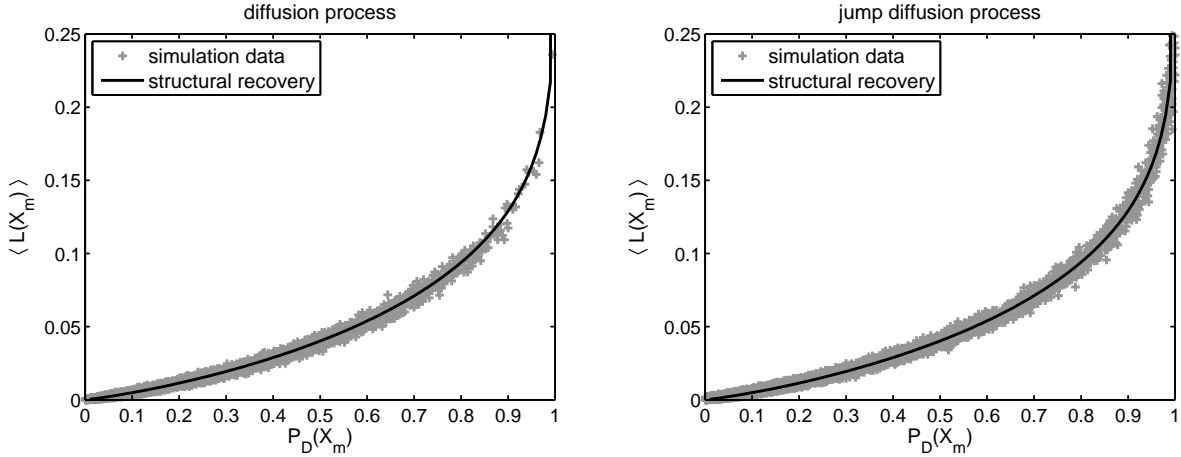


Figure 2: Dependence of average losses $\langle L(X_m) \rangle$ on default probabilities $P_D(X_m)$. Results are shown for diffusion (left), and for jump–diffusion (right).

5. Comparison of model results

A. Dependence of recovery rates and default probabilities

Let us first turn our attention to the dependence between recovery rates and default probabilities. Figure 3 shows scatter plots of the Monte Carlo results for $P_D(X_m)$ and $\langle R(X_m) \rangle$. We compare these with the results for the three different recovery rate models, where $P_D(X_m)$ is again taken directly from the simulation. Here, the model parameters for the recovery rate are calibrated to the data with $X_m < 0$. It is obvious that the assumption of a constant recovery rate is not justified at all. Moreover, the estimation of this constant is biased due to the relative abundance of data for small default probability and high recovery rate. Although this bias could be taken into account, the results for the constant recovery rate model appear arbitrary and neglect important structural information. For the reduced–form recovery, we notice that the interdependence of recovery and default rates is preserved at least qualitatively. We observe larger deviations for the jump–diffusion. This is to be expected, since the recovery model does not describe the simulation results over the whole data range, see our discussion in the previous section. The qualitative agreement of the reduced–form model results and the simulation data can be attributed to the fact that both $P_D(X_m)$ and $\langle R(X_m) \rangle$ describe the dependence on the same covariate. This is in

accordance with findings of Chava et al. (2008) who pointed out that reduced-form models can yield considerably better results, if defaults and recoveries are modelled and calibrated with respect to the same set of covariates.

The structural recovery describes the data very well over the whole data range. This is not surprising for the diffusion process, after all the functional dependence has been derived for this case. It is noteworthy, however, that we find such a perfect agreement for the jump-diffusion as well. This result was previously mentioned in Schäfer and Koivusalo (2011).

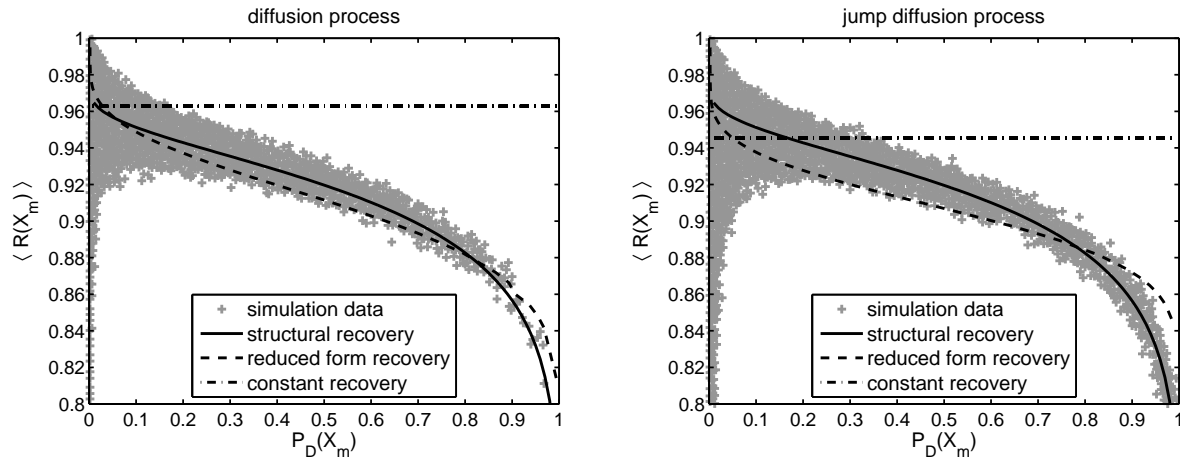


Figure 3: Dependence of recovery rates $\langle R(X_m) \rangle$ on default probability $P_D(X_m)$. The Monte Carlo results are compared to structural (solid line), reduced-form (dashed line) and constant recovery rate (dashed-dotted line). Results are shown for diffusion (left), and for jump-diffusion (right).

B. Value at Risk and Expected Tail Loss

Finally, we want to test the quality of the model results on two risk measures that capture the tail behavior of the loss distribution. For this purpose, we consider the Value at Risk and the Expected Tail Loss, or Expected Shortfall. Both measures are commonly used in credit risk management, see, e.g., Artzner et al. (1997); Frey and McNeil (2002); Yamai and Yoshida (2005), and in financial regulations, see Basel Committee on Banking Supervision (2006). For details on the calculation of these risk measures we refer to Appendix A. As mentioned above, we are considering a situation where the default probabilities are well described by a model — or in our case by the simulation itself. Hence, we focus entirely on the influence of the different recovery rate models. In particular, we are interested in the impact of limited calibration data on the model results. To this end we filter the simulation data with respect to the market return X_m . As before we keep 0 as an upper threshold, because default and recovery rates do not show a clear dependence on the market return for positive X_m . Additionally, we set a lower threshold on X_m and study the results for VaR and ETL in dependence of this lower threshold. This corresponds to the relative abundance of empirical data for moderate market returns, while large negative market returns occur only scarcely. Frankly, we ask to which extent the different models can describe extreme

events when they are only calibrated with default and recovery data from a relatively calm period.

Figure 4 shows the model results for the Value at Risk VaR_α at the confidence level $\alpha = 0.99$. The results are normalized to the corresponding empirical value which is computed for the entire data set of the simulation. These empirical values are $\text{VaR}_{0.99} = 0.013$ for the diffusion and $\text{VaR}_{0.99} = 0.015$ for the jump-diffusion. Not surprisingly, we find the worst results for the constant recovery. The Value at Risk is underestimated by more than 30% in the diffusion case and by almost 20% in the jump-diffusion case. The better result for the jump-diffusion is misleading, however, as we will see later on. The reduced-form recovery yields quite reasonable results in the diffusion case. It overestimates the Value at Risk by roughly 8% when all data with $X_m < 0$ is taken into account. We see a slight dependence on the lower threshold for X_m . This starts only at $X_m \approx -0.35$, because this was the lowest market return observed in the simulation. In the jump-diffusion case, the reduced-form model yields very unstable results, ranging from a considerable overestimation to a severe underestimation of the Value at Risk. For both processes, the structural recovery model provides rather stable results with very little deviation from the empirical values.

In Figure 5 we present the model results for the Expected Tail Loss ETL_α at the confidence level $\alpha = 0.99$. Again, we normalized the results to the empirical value for the entire data set of the simulation. These empirical values are $\text{ETL}_{0.99} = 0.0238$ for the diffusion and $\text{ETL}_{0.99} = 0.0379$ for the jump-diffusion. The Expected Tail Loss is better suited to describe the tail behavior of the loss distribution, as first pointed out by Artzner et al. (1997). Here we see that the constant recovery rate does in fact not provide a better description of the jump-diffusion than of the diffusion. The results for the Value at Risk are more coincidental, they depend critically on the chosen confidence level α . The reduced-form recovery again shows rather good results for the diffusion process. However, the Expected Tail Loss reveals even more clearly that the reduced-form recovery severely underestimates large losses, if the calibration data is restricted. Moreover, we observe a very strong dependence of the results on the lower threshold for the market returns. These unstable results are contrasted by the stable results for the structural recovery rate. The Expected Tail Loss confirms that this model delivers a very good and robust estimate for large portfolio losses.

6. Conclusions

The relation between recovery rates and default probabilities has a pronounced effect on the tail of the portfolio loss distribution. Yet both quantities are often modelled independently in current credit risk models. In this paper we consider a Merton model as a reference system with diffusion and jump-diffusion as underlying processes. Three different recovery models are calibrated to Monte Carlo simulations of this reference system: a reduced-form recovery model, the special case of constant recoveries, and the structural recovery model. A functional dependence between default and recovery rates can be derived for the Merton model with correlated diffusion. The structural recovery model employs this same functional dependence regardless of the underlying process. We study to what extent the different recovery models preserve the relation between default and recovery rates. Most

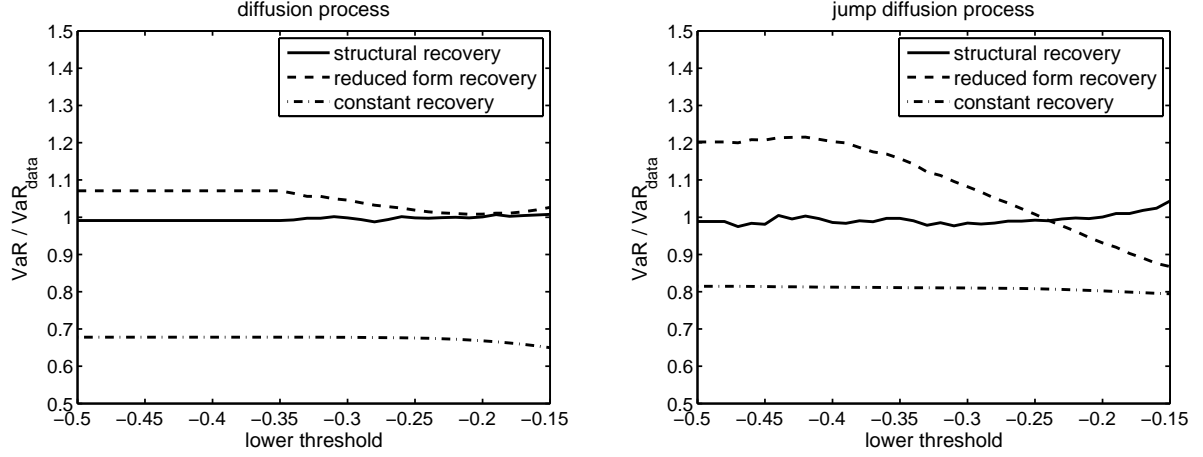


Figure 4: Value at Risk VaR_α at the confidence level $\alpha = 0.99$. The results are normalized to the Value at Risk computed for the entire data set of the simulations. Thus, values above 1 correspond to an overestimation, values below 1 to an underestimation by the model. Results are shown for diffusion (left), and for jump-diffusion (right).

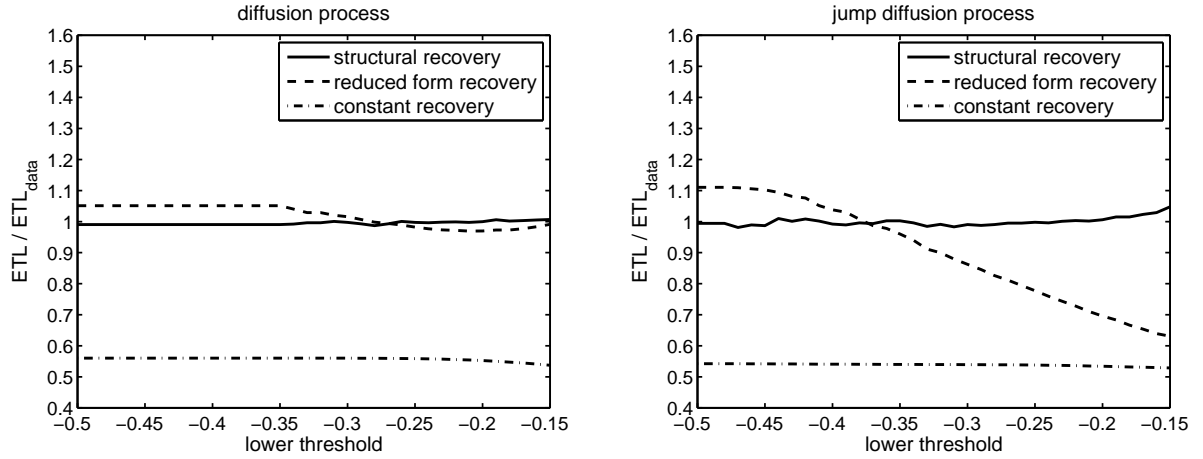


Figure 5: Expected Tail Loss ETL_α at the confidence level $\alpha = 0.99$. The results are normalized to the ETL computed for the entire data set of the simulations. Thus, values above 1 correspond to an overestimation, values below 1 to an underestimation by the model. Results are shown for diffusion (left), and for jump-diffusion (right).

importantly, we compare how well they reproduce two measures of tail risk, the Value at Risk and the Expected Tail Loss. Furthermore, we discuss calibration issues that may arise due to a lack of information. To this end we confine the calibration to a certain range of the simulation data. A lower threshold on the market return implies that extremely adverse market situations are not contained in the historical data sample on defaults and recoveries.

It comes as no surprise that the constant recovery rate model shows the worst performance, since it does not include a dependence on covariates or default probabilities. In addition, the estimate of the average recovery rate from historical or simulation data is bi-

ased due to the relative abundance of small losses. As a consequence this model shows very poor results for Value at Risk and Expected Tail Loss. Nonetheless we included this model in our discussion, as it is still prevalent in practical applications.

The results for the reduced-form recovery model are considerably better. The generalized linear dependence on market returns works reasonably well, especially when considering large negative market returns. However, it does not provide a good description over entire data range. Furthermore, the quality of the model critically depends on the underlying process. The model is able to reproduce the dependence of default and recovery rates at least qualitatively. Larger deviations are observed for the jump diffusion process. We find a critical dependence of the model results on the data range used for calibration. This leads to very unstable estimates for Value at Risk and Expected Tail Loss.

The structural recovery model accurately describes the dependence of default and recovery rates for different processes. It has only a single parameter and is calibrated with respect to the default rate dependence. The model yields nearly perfect estimates of Value at Risk and Expected Tail Loss. These results show almost no dependence on the data range used for calibration.

In accordance with Chava et al. (2008) we conclude that a reduced-form ansatz can work well, if the same covariates are used for both default and recover rates. However, an additional model risk is involved in the reduced-form recovery with respect to the underlying process. The structural recovery avoids this model risk because it does not rely on covariates, instead it directly describes the dependence of recovery rates on default probabilities. Thus, it can be easily used with any default model.

Acknowledgments

We wish to thank Thomas Guhr and Sven Åberg for helpful discussions.

Appendix A. Calculation of risk measures

In the reduced-form approach default probabilities and recovery rates are modelled in dependence of covariates, in our case, in dependence of the market return X_m . The average loss or loss of a homogeneous portfolio is then

$$\langle L(X_m) \rangle = P_D(X_m) (1 - \langle R(X_m) \rangle) . \quad (\text{A.1})$$

There is an abundance of empirical data on covariates, we can therefore assume that their distribution is well understood and modelled. In the case of a correlated diffusion, we know that the market return is log-normal distributed. The jump-diffusion process renders the tail of the distribution fatter, corresponding more closely to the situation on the stock market. When the PDF $p_{X_m}(X_m)$ is known, we can use the substitution

$$|p_L(L)dL| = |p_{X_m}(X_m)dX_m| \quad (\text{A.2})$$

to calculate the Value at Risk and the Expected Tail Loss. Here we simplified the notation to L instead of $\langle L(X_m) \rangle$. With the cumulative distribution function

$$F_{X_m}(x) = \int_{-\infty}^x p_{X_m}(X_m) dX_m \quad (\text{A.3})$$

we can express the Value at Risk VaR_α as

$$\text{VaR}_\alpha = L(F_{X_m}^{-1}(1 - \alpha)) , \quad (\text{A.4})$$

where the function L is the X_m -dependence given in Equation (A.1). The value $F_{X_m}^{-1}(1 - \alpha)$ is the $1 - \alpha$ quantile of the market return. For numerical data it can be calculated directly.

The Expected Tail Loss is calculated as

$$\text{ETL}_\alpha = \frac{1}{\alpha} \int_{\text{VaR}_\alpha}^1 L p_L(L) dL = \frac{1}{\alpha} \int_{-\infty}^{F_{X_m}^{-1}(1-\alpha)} L(X_m) p_{X_m}(X_m) dX_m . \quad (\text{A.5})$$

For numerical data we can determine the Expected Tail Loss as

$$\text{ETL}_\alpha = \text{mean} \{ L(X_m) \mid X_m \leq F_{X_m}^{-1}(1 - \alpha) \} . \quad (\text{A.6})$$

The same line of reasoning applies for the structural recovery approach. The average loss is then given by

$$\langle L(P_D) \rangle = P_D (1 - \langle R(P_D) \rangle) , \quad (\text{A.7})$$

with the expected recovery rate according to Equation (13). When we know the distribution of default probabilities, e.g., by means of modelling, by a scenario analysis or from Monte Carlo simulations, we can express the Value at Risk as

$$\text{VaR}_\alpha = L(F_{P_D}^{-1}(\alpha)) , \quad (\text{A.8})$$

where F_{P_D} is the cumulative distribution function of the default probability. For the Expected Tail Loss we find

$$\text{ETL}_\alpha = \frac{1}{\alpha} \int_{F_{P_D}^{-1}(\alpha)}^1 L(P_D) p_{P_D}(P_D) dP_D . \quad (\text{A.9})$$

For numerical data we can evaluate it as

$$\text{ETL}_\alpha = \text{mean} \{ L(P_D) \mid P_D \geq F_{P_D}^{-1}(\alpha) \} . \quad (\text{A.10})$$

References

- Andersen, L., Sidenius, J., 2005. Extensions to the gaussian copula: random recovery and random factor loadings. *Journal of Credit Risk* 1 (1), 29–70.
- Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1997. Thinking coherently. *Risk* 10, 68–71.
- Asvanunt, A., Staal, A., April 2009a. The Corporate Default Probability model in Barclays Capital POINT platform (POINT CDP). Portfolio Modeling, Barclays Capital.
- Asvanunt, A., Staal, A., August 2009b. The POINT Conditional Recovery Rate (CRR) Model. Portfolio Modeling, Barclays Capital.
- Basel Committee on Banking Supervision, June 2006. International Convergence of Capital Measurement and Capital Standards. BIS, <http://www.bis.org/publ/bcbs128.htm>.
- Black, F., Cox, J. C., 1976. Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance* 31, 351–367.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81 (3), 637.
URL <http://www.journals.uchicago.edu/doi/abs/10.1086/260062>
- Chava, S., Stefanescu, C., Turnbull, S., 2008. Modeling the loss distribution, working paper.
- Duffie, D., Singleton, K., 1999. Modeling the term structure of defaultable bonds. *Review of Financial Studies* 12, 687–720.
- Frey, R., McNeil, A. J., 2002. Var and expected shortfall in portfolios of dependent credit risks: Conceptual and practical insights. *Journal of Banking & Finance* 26 (7), 1317–1334.
- Giesecke, K., 2004. Credit Risk Modeling and Valuation: An Introduction, 2nd Edition. Credit Risk: Models and Management. Risk Books, Ch. 16, p. 487.
- Hull, J. C., White, A., 2000. Valuing credit default swaps I: No counterparty default risk. *Journal of Derivatives* 8 (1), 29–40.
- Jarrow, R. A., Lando, D., Turnbull, S. M., 1997. A markov model for the term structure of credit risk spreads. *Review of Financial Studies* 10 (2), 481–523.
- Jarrow, R. A., Turnbull, S. M., 1995. Pricing derivatives on financial securities subject to default risk. *Journal of Finance* 50, 53–86.
- Kiesel, R., Scherer, M., 2007. Dynamic credit portfolio modelling in structural models with jumps, preprint available at http://www.defaultrisk.com/pp_model170.htm.
- Mantegna, R., Stanley, H., 2000. An Introduction to Econophysics. Cambridge University Press, Cambridge.
- Merton, R. C., 1974. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance* 29, 449–470.
- Schäfer, R., Koivusalo, A., 2011. Dependence of defaults and recoveries in structural credit risk models. Preprint, arXiv:1102.3150.
- Schäfer, R., Sjölin, M., Sundin, A., Wolanski, M., Guhr, T., 2007. Credit risk — a structural model with jumps and correlations. *Physica A* 383 (2), 533.
- Schönbucher, P. J., 2003. Credit Derivatives Pricing Models. John Wiley & Sons, New Jersey.
- Yamai, Y., Yoshida, T., 2005. Value-at-risk versus expected shortfall: A practical perspective. *Journal of Banking & Finance* 29 (4), 997–1015.
- Zhou, C., 2001. The term structure of credit spreads with jump risk. *Journal of Banking and Finance* 25, 2015.